

A Reduction Theorem for the Sample Mean in Dynamic Time Warping Spaces

Brijnesh Jain and David Schultz
 Technische Universität Berlin, Germany
 e-mail: brijnesh.jain@gmail.com

Abstract. Though the concept of sample mean in dynamic time warping (DTW) spaces is used in pattern recognition applications, its existence has neither been proved nor called into question. This article shows that a sample mean exists under general conditions that cover common variations of different DTW-spaces mentioned in the literature. The existence proofs are based on a Reduction Theorem that bounds the length of the candidate solutions we need to consider. The proposed results place the concept of sample mean in DTW-spaces on a sound mathematical foundation and serves as a first step towards a statistical theory of DTW-spaces.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 2 | The Reduction Theorem and its Implications | 2 |
| 2.1 | The Dynamic Time Warping Distance | 2 |
| 2.2 | Fréchet Functions | 4 |
| 2.3 | The Reduction Theorem | 4 |
| 2.4 | Sufficient Conditions of Existence | 5 |
| 3 | Theory of Warping Graphs | 8 |
| 3.1 | Warping Chains | 8 |
| 3.2 | Warping Graphs | 9 |
| 3.3 | Glued Warping Graphs | 13 |
| 3.4 | Labeled Warping Graphs | 16 |
| 3.5 | Proofs of Results from Section 2 | 16 |
| 4 | Conclusion | 19 |

1 Introduction

The sample mean in Euclidean spaces is the epitome of a summary measure that attempts to describe a whole set of data with a single value. In addition, the sample mean is one of the most fundamental concepts in statistical inference and also forms the basis of pattern recognition methods such as, for example, Gaussian mixture models and principal component analysis. The sample mean is theoretically well understood in Euclidean spaces, but can easily become obscure in mathematically less structured spaces. Since an increasing amount of non-Euclidean data is being collected and analyzed in ways that have not been realized before, statistics is undergoing an evolution [12]. Evidence of this evolution are contributions to statistical analysis of shapes [2, 3, 9, 11], complex objects [14, 20], tree-structured data [4, 20], and graphs [6, 10].

All the above mentioned spaces have in common that they are metric spaces for which a well-defined addition is unknown. In 1948, Fréchet [5] generalized the idea of averaging in arbitrary metric spaces by using a sum-of-distances criterion. Following Fréchet’s approach, a sample mean of a sample $\mathcal{X} = (x^{(1)}, \dots, x^{(N)})$ of N time series $x^{(k)}$ is any time series that minimizes the Fréchet function [5]

$$F(x) = \sum_{k=1}^N \delta(x, x^{(k)}), \quad (1)$$

where $\delta(x, y)$ is the dynamic time warping (DTW) distance [18]. Fréchet’s idea has been adopted to DTW-spaces in early work by Rabiner & Wilpon [17] in the late 1970ies. Since then, research on the sample mean predominantly focuses on devising algorithms for approximating a sample mean of time series [1, 7, 8, 13, 15, 16, 19].

One of the most basic questions of any optimization problem and of statistical inference remained unanswered for nearly four decades, namely the *existence of a sample mean in DTW-spaces*. Without any evidence of existence, algorithms have been proposed for finding solutions that potentially may not exist. Moreover, the approximations returned by the algorithms have been used as prototypes in nearest neighbor classification or k-means clustering. Though this approach may yield satisfactory results in applications, it lacks a theoretical justification. In addition, the lack of a sound theory of the sample mean proliferates in branches of statistical inference and pattern recognition.

This article presents a Reduction Theorem, which considerably simplifies existence proofs and points to a technique for improving candidate solutions. Based on the Reduction Theorem, we show existence of a sample mean under general conditions that cover common DTW-spaces mentioned in the literature.

The line of argument to prove existence of a sample mean is as follows: First, we consider two forms of the Fréchet function, a restricted and an unrestricted form. The restricted Fréchet function is defined on the subset of all time series of fixed length m . In contrast, the unrestricted Fréchet function imposes no restrictions on the length of the time series. In either of both forms, the sample time series $x^{(k)}$ can have any length. Then we show existence of a restricted sample mean for every length. Finally, we infer existence of an unrestricted sample mean. The Reduction Theorem enables the transition from restricted to unrestricted sample means. The proposed theorem states that there is a sample-dependent bound on the length beyond which the Fréchet function can not be improved. The proof of the Reduction Theorem is framed into a novel theory of warping graphs.

Existence of restricted sample means theoretically justify prior work [1, 8, 16, 17] and their follow-up papers. The unrestricted form advances research on the sample mean in DTW-spaces by showing that the study of sample means can be reduced to the study of restricted sample means of bounded length. The proposed results place the concept of sample mean on a sound mathematical foundation and serves as a first step towards a statistical theory of DTW-spaces.

This rest of this paper is structured as follows: Section 2 states the main results of this contribution. Section 3 first develops a theory of warping graphs and then proves the results of Section 2. Finally, Section 4 concludes with a summary of the main findings and an outlook to further research.

2 The Reduction Theorem and its Implications

This section first introduces the DTW-distance and Fréchet functions. Then the Reduction Theorem is stated and its implications are presented. Finally, sufficient conditions of existence of a sample mean are proposed.

Notations. We write $\mathbb{R}_{\geq 0}$ for the set of non-negative reals. By \mathbb{N} we denote the set of positive integers. We write $[n]$ to denote the set $\{1, \dots, n\}$ for a given $n \in \mathbb{N}$. Let \mathcal{S} be a set. Then $\mathcal{S}^N = \mathcal{S} \times \dots \times \mathcal{S}$ is the N -fold Cartesian product of \mathcal{S} , where $N \in \mathbb{N}$.

2.1 The Dynamic Time Warping Distance

Without further mention, we assume that \mathcal{A} is an attribute set and $d : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}$ is a non-negative distance function on \mathcal{A} , called *local cost function*. A *time series* x of *length* $\ell(x) = m$ is an ordered

sequence $x = (x_1, \dots, x_m)$ consisting of elements $x_i \in \mathcal{A}$ for every time point $i \in [m]$. By \mathcal{T}_n we denote the set of all time series of length $n \in \mathbb{N}$ with elements from \mathcal{A} . Then

$$\mathcal{T} = \bigcup_{n \in \mathbb{N}} \mathcal{T}_n$$

is the set of all time series of finite length with elements from \mathcal{A} . The DTW-distance is a distance function on \mathcal{T} based on the notion of warping path.

Definition 2.1. Let $m, n \in \mathbb{N}$. A warping path of order $m \times n$ is a sequence $p = (p_1, \dots, p_L)$ of L points $p_l = (i_l, j_l) \in [m] \times [n]$ such that

1. $p_1 = (1, 1)$ and $p_L = (m, n)$ (boundary conditions)
2. $p_{l+1} - p_l \in \{(1, 0), (0, 1), (1, 1)\}$ for all $l \in [L - 1]$ (step condition)

The set of all warping paths of order $m \times n$ is denoted by $\mathcal{P}_{m,n}$. A warping path of order $m \times n$ can be thought of as a path in a $[m] \times [n]$ grid, where rows are ordered top-down and columns are ordered left-right. The boundary condition demands that the path starts at the upper left corner and ends in the lower right corner of the grid. The step condition demands that a transition from one point to the next point moves a unit in exactly one of the following directions: down, diagonal, and right.

A warping path $p = (p_1, \dots, p_L) \in \mathcal{P}_{m,n}$ defines an alignment (or warping) between time series $x = (x_1, \dots, x_m)$ and $y = (y_1, \dots, y_n)$. Every point $p_l = (i_l, j_l)$ of warping path p aligns element x_{i_l} to element y_{j_l} . The cost of aligning time series x and y along warping path p is defined by

$$c_p(x, y) = \sum_{l=1}^L d(x_{i_l}, y_{j_l}),$$

where d is the aforementioned local distance function on \mathcal{A} . Then the DTW-distance between two time series minimizes the cost of aligning both time series over all possible warping paths.

Definition 2.2. Let $f : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ be a monotonous function. Let x and y be two time series of length m and n , respectively. The DTW-distance between x and y is defined by

$$\delta(x, y) = \min \{f(c_p(x, y)) : p \in \mathcal{P}_{m,n}\}.$$

An optimal warping path is any warping path $p \in \mathcal{P}_{m,n}$ satisfying $\delta(x, y) = f(c_p(x, y))$.

Even if the underlying local cost function is a metric, the induced DTW-distance is generally only a pseudo-semi-metric satisfying

1. $\delta(x, y) \geq 0$
2. $\delta(x, x) = 0$

for all $x, y \in \mathcal{T}$. Computing the DTW-distance and deriving an optimal warping path is usually solved by applying techniques from dynamic programming [18]. The next example presents a common and widely applied DTW-distance.

Example 2.3. The Euclidean DTW-distance is specified by the attribute set $\mathcal{A} = \mathbb{R}^d$, the squared Euclidean distance $d(x, y) = \|x - y\|^2$ for all $x, y \in \mathcal{A}$, and the square root function $f(x) = \sqrt{x}$ for all $x \in \mathbb{R}_{\geq 0}$.

A DTW-space is a pair (\mathcal{T}, δ) consisting of a set of time series of finite length and a DTW-distance δ defined on \mathcal{T} . For the sake of convenience, we occasionally write \mathcal{T} to denote a DTW-space and tacitly assume that δ is the underlying DTW-distance.

2.2 Fréchet Functions

We assume that (\mathcal{T}, δ) is a DTW-space. To admit different measures of central location such as the sample mean and median in the case of Euclidean spaces, we consider a more general form of Fréchet function than the sum-of-distances criterion introduced in Section 1, Eq. (1).

A *loss function* is a monotonously increasing function of the form $h : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$. A typical example of a loss function is the squared loss $h(u) = u^2$ for all $u \geq 0$.

Definition 2.4. Let $h : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ be a loss function. Suppose that $\mathcal{X} = (x^{(1)}, \dots, x^{(N)}) \in \mathcal{T}^N$ is a sample of N time series. Then the function

$$F : \mathcal{T} \rightarrow \mathbb{R}, \quad x \mapsto \frac{1}{N} \sum_{i=1}^N h\left(\delta\left(x, x^{(k)}\right)\right)$$

is the Fréchet function of sample \mathcal{X} corresponding to the loss h .

We omit explicitly mentioning the corresponding loss function of a Fréchet function if no confusion can arise.

Example 2.5. Let $p \geq 1$. Consider the loss function $h(u) = u^p$ for all $u \in \mathbb{R}_{\geq 0}$. Then the Fréchet function of \mathcal{X} corresponding to the loss $h(u) = u^p$ takes the form

$$F^p(x) = \frac{1}{N} \sum_{i=1}^N \delta^p\left(x, x^{(k)}\right).$$

For $p = 1$, the Fréchet function F^p coincides with the sum-of-distances criterion as defined in Eq. (1). In addition, for $p = 1$ ($p = 2$) the Fréchet function F^p serves to generalize the concept of sample median (sample mean) in Euclidean spaces.

Definition 2.6. The sample mean set of \mathcal{X} is the (possibly empty) set defined by

$$\mathcal{F} = \{z \in \mathcal{T} : F(z) \leq F(x) \text{ for all } x \in \mathcal{T}\}.$$

The elements of \mathcal{F} are the sample means of \mathcal{X} .

A sample mean is a time series that minimizes the Fréchet function F . Existence of a sample mean depends on the choice of DTW-distance and loss function. Moreover, if a sample mean exists, it may not be uniquely determined. In contrast, the *sample variation*¹

$$F^* = \inf_{x \in \mathcal{T}} F(x)$$

exists and is uniquely determined, because the DTW-distance is bounded from below and the loss is monotonously increasing. Thus, existence of a sample mean means that the Fréchet function F attains its infimum.

2.3 The Reduction Theorem

Suppose that $\mathcal{X} = (x^{(1)}, \dots, x^{(N)}) \in \mathcal{T}^N$ is a sample of N time series. The Reduction Theorem is based on the notion of *reduction bound* $\rho(\mathcal{X})$ of sample \mathcal{X} . The exact definition of $\rho(\mathcal{X})$ requires some technicalities and is fully spelled out in Section 3.3. Here, we present a simpler definition that conveys the main idea and covers the relevant use cases in pattern recognition. For this, we assume that every sample time series $x^{(k)}$ of sample \mathcal{X} has length $\ell(x^{(k)}) \geq 2$. Then the reduction bound of \mathcal{X} is defined by

$$\rho(\mathcal{X}) = \sum_{k=1}^N \ell\left(x^{(k)}\right) - 2(N-1). \quad (2)$$

¹The sample variation F^* corresponding to loss $h(u) = u^2$ generalizes the sample variance in Euclidean spaces.

In contrast to Eq. (2), the exact definition of $\rho(\mathcal{X})$ admits samples that contain trivial time series of length one. Equation (2) shows that the reduction bound of a sample increases linearly with the sum of the lengths of the sample time series. The next as well as subsequent results hold for arbitrary samples and the exact definition of a reduction bound as provided in Section 3.3.

Theorem 2.7 (Reduction Theorem). *Let F be the Fréchet function of a sample $\mathcal{X} \in \mathcal{T}^N$. Then for every time series $x \in \mathcal{T}$ of length $\ell(x) > \rho(\mathcal{X})$ there is a time series $x' \in \mathcal{T}$ of length $\ell(x') = \ell(x) - 1$ such that $F(x') \leq F(x)$.*

The Reduction Theorem deserves some explanations. To illustrate the following comments we refer to Figures 1 – 3. These figures assume univariate time series with continuous values. The underlying distance is the Euclidean DTW-distance of Example 2.3. The Fréchet functions of the different samples $\mathcal{X} = (x^{(1)}, x^{(2)})$ correspond to the identity loss $h(u) = u$ and thus coincide with the function F^p of Example 2.5, where $p = 1$. The figures show warping paths by black lines connecting aligned elements of the time series to be compared. We make the following observations:

1. From the proof of the Reduction Theorem follows that every time series x whose length exceeds the reduction bound has an element that can be removed without increasing the value $F(x)$. Such elements are said to be *redundant*. Figure 1 schematically characterizes redundant elements of a time series.
2. In general, removing a redundant element does not increase the Fréchet function. Figure 1 shows that removing a redundant element can even decrease the value of the Fréchet function.
3. The reduction bound of the sample in Figure 1 is given by

$$\rho(\mathcal{X}) = \ell(x^{(1)}) + \ell(x^{(2)}) - 2(N - 1) = 4 + 4 - 2 = 6.$$

The length of time series x is only $\ell(x) = 5 < \rho(\mathcal{X})$. This shows that short time series whose lengths are bounded by the reduction bound may also have redundant elements that can be removed without increasing the value of the Fréchet function. Existence of a redundant element depends on the choice of warping paths between x and the sample time series. For short time series x , we can always find warping paths such that x has no redundant elements. In contrast, long time series whose lengths exceed the reduction bound always have a redundant element, regardless which warping paths we consider.

4. Removing a non-redundant element of a time series can increase the value of the Fréchet function. Figure 2 presents an example.
5. The Reduction Theorem does not exclude existence of sample means whose lengths exceed the reduction bound of a sample. Figure 3 presents an example that a sample mean can have almost any length.

The Reduction Theorem and observations 1–4 form the basis for existence proofs and point to a technique to improve algorithms for approximating a sample mean. Statements on the existence of a sample mean are presented in the next section. From observations 1–3 follows that a candidate solution x of any length could be improved by detecting and removing redundant elements of x . This observation is not further explored in this article and left for further research.

2.4 Sufficient Conditions of Existence

In this section, we derive sufficient conditions of existence of a sample mean in DTW-spaces.

The Reduction Theorem is useful, because it considerably simplifies existence proofs. The general approach is to prove existence of a sample mean for a class of restricted problems and then infer existence of a sample mean for the unrestricted case by using the Reduction Theorem.




Figure 1: Schematic description of redundant elements. Shown are two sample time series $x^{(1)}$ and $x^{(2)}$ and two further time series x and x' . The third element of x is redundant and therefore highlighted by an enclosing circle. Redundant elements are characterized by the following property: An element of time series x is redundant if every element of the sample time series connected to x is connected to at least two elements of x . The elements of the sample time series connected to the third element of x are enclosed by a square. Both squared elements are connected to two elements of x . The time series x' is obtained from x by removing its redundant element. The DTW-distances of the sample time series from x are both one and from x' are both zero. Then we obtain $F(x) > F(x')$. This shows that removing the redundant element in x does not increase the value of the Fréchet function. In this particular case, the value of the Fréchet function is even decreased.

To describe the class of restricted problems, we assume that $\mathcal{X} = (x^{(1)}, \dots, x^{(k)}) \in \mathcal{T}^N$ is a sample and h is a loss function. Then the function

$$F_m : \mathcal{T}_m \rightarrow \mathbb{R}, \quad x \mapsto F(x).$$

is the Fréchet function of \mathcal{X} restricted to the subset $\mathcal{T}_m \subset \mathcal{T}$ of all time series of length m . Note that the lengths of the sample time series in \mathcal{X} may vary, but the length of the independent variable x of $F_m(x)$ is fixed beforehand to value m . The set

$$\mathcal{F}_m = \{z \in \mathcal{T}_m : F_m(z) \leq F_m(x) \text{ for all } x \in \mathcal{T}_m\}$$

is the *restricted sample mean set* of \mathcal{X} restricted to the subset \mathcal{T}_m . We show that existence of a class of restricted sample means implies existence of a sample mean.

Corollary 2.8. *Let $\mathcal{X} \in \mathcal{T}^N$ be a sample and let $\rho \in \mathbb{N}$ be the reduction bound of \mathcal{X} . Suppose that $\mathcal{F}_m \neq \emptyset$ for every $m \in [\rho]$. Then \mathcal{X} has a sample mean.*

It is not self-evident that existence of a class of restricted sample mean implies existence of a sample mean. To see this, we define the restricted (sample) variation by

$$F_m^* = \inf_{x \in \mathcal{T}_m} F_m(x).$$

If F_m attains its infimum, then \mathcal{X} has a restricted sample mean. Suppose that \mathcal{X} has a restricted sample mean for every $m \in \mathbb{N}$. Then

$$v_m = \min_{l \leq m} F_l^*.$$

is the smallest restricted variation over all lengths $l \in [m]$. The sequence $(v_m)_{m \in \mathbb{N}}$ is bounded from below and monotonously decreasing. Therefore, the $(v_m)_{m \in \mathbb{N}}$ converges to the unrestricted sample variation F_* . Then \mathcal{X} has a sample mean only if the sequence $(v_m)_{m \in \mathbb{N}}$ attains its infimum F_* . Corollary 2.8 guarantees that the sequence $(v_m)_{m \in \mathbb{N}}$ attains its infimum F_* latest at $m = \rho(\mathcal{X})$.




Figure 2: Removing a non-redundant element can increase the Fréchet function. The time series x' is obtained from x by removing the first element. The first element of x is not redundant according to the characterization of redundant elements given in Figure 1. The DTW-distances of the sample time series from x are both zero and from x' are both one. This shows that $F(x) < F(x')$.




Figure 3: Sample means of almost any length. The time series x is a sample mean of \mathcal{X} , because $F(x) = 0$ is the lowest possible value. The time series x' is obtained from the sample mean x by appending arbitrarily many time points with element 3. From $F(x') = 0$ follows that x' is also a sample mean of \mathcal{X} .

We conclude this section with presenting sufficient conditions of existence. The first result proposes sufficient conditions of existence of a (restricted) sample mean for time series (sequences) with discrete attribute values.

Proposition 2.9. *Let $\mathcal{X} \in \mathcal{T}^N$ be a sample. Suppose that \mathcal{A} is a finite attribute set. Then the following statements hold:*

1. $\mathcal{F}_m \neq \emptyset$ for every $m \in \mathbb{N}$.
2. $\mathcal{F} \neq \emptyset$.

The second result proposes sufficient conditions of existence of a (restricted) sample mean of uni- and multivariate time series with elements from $\mathcal{A} = \mathbb{R}^d$.

Proposition 2.10. *Let $\mathcal{X} \in \mathcal{T}^N$ be a sample. Suppose that the following assumptions hold:*

1. (\mathcal{A}, d) is a metric space of the form $(\mathbb{R}^d, \|\cdot\|)$, where $\|\cdot\|$ is a norm on \mathbb{R}^d .
2. The loss h is continuous and strictly monotonously increasing.

Then the following statements hold:

1. $\mathcal{F}_m \neq \emptyset$ for every $m \in \mathbb{N}$.
2. $\mathcal{F} \neq \emptyset$.

The attribute set \mathcal{A} covers the case of univariate ($d = 1$) and the case of multivariate ($d > 1$) time series. The local cost function d on \mathcal{A} is a metric induced by a norm on \mathbb{R}^d . The loss function $h : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ with $h(u) = u^p$ is continuous and strictly monotonously increasing for a fixed $p \geq 1$. Thus, the sufficient conditions of Prop. 2.10 cover customary DTW-spaces.

3 Theory of Warping Graphs

This section develops a theory of warping graphs to prove the Reduction Theorem. The line of argument follows a bottom-up approach. First, we introduce warping chains to model abstract warping paths as given in Definition 2.1 and derive their relevant local properties. Then we proceed to warping graphs that model the alignment of two time series by a warping path and derive global properties from local properties. We enhance warping graphs with node labels and define the notion of weight of a warping graph to model the DTW-distance. Finally, we glue warping graphs to model Fréchet functions, derive the Reduction Theorem, and prove the statements presented in Section 2.

3.1 Warping Chains

The basic constituents of a Fréchet function are time series and optimal warping paths. This section represents the linear order of time series by chains. Then we introduce the notion of warping chain to model abstract warping paths and study its local properties.

Let \mathcal{V} be a partially ordered set with partial order $\leq_{\mathcal{V}}$. Suppose that $i, j \in \mathcal{V}$ are two elements with $i \leq_{\mathcal{V}} j$. We write $i <_{\mathcal{V}} j$ to mean $i \leq_{\mathcal{V}} j$ and $i \neq j$. A linear order $\leq_{\mathcal{V}}$ on \mathcal{V} is a partial order such that any two elements in \mathcal{V} are comparable: For all $i, j \in \mathcal{V}$, we have either $i \leq_{\mathcal{V}} j$ or $j \leq_{\mathcal{V}} i$.

Definition 3.1. *A chain is a linearly ordered set.*

A chain \mathcal{V} models the order of a time series $x = (x_1, \dots, x_m)$, where element $i \in \mathcal{V}$ refers to the positions of element x_i in x . For the sake of convenience, we assume that the explicit notation of a chain $\mathcal{V} = \{i_1, \dots, i_m\}$ always lists its elements in linear order $i_1 <_{\mathcal{V}} \dots <_{\mathcal{V}} i_m$. We call i_1 the first and i_m the last element in \mathcal{V} . The first and last element of a chain are the boundary elements. Any element of chain \mathcal{V} that is not a boundary element is called an inner element of \mathcal{V} .

A subset $\mathcal{V}' \subseteq \mathcal{V}$ is a subchain of \mathcal{V} . Note that any subset of a chain is again a chain by transitivity of the linear order. Suppose that $i_p, i_q \in \mathcal{V}$ such that $i_p \leq_{\mathcal{V}} i_q$. Then the subchain $\mathcal{V}' = \{i \in \mathcal{V} : i_p \leq_{\mathcal{V}} i \leq_{\mathcal{V}} i_q\}$ is said to be contiguous.

Let $\mathcal{V} = \{i_1, \dots, i_m\}$ be a chain and let $\mathcal{V}^* = \mathcal{V} \cup \{*\}$, where $*$ is a distinguished symbol denoting the void element. The successor i_l^+ and predecessor i_l^- of element $i_l \in \mathcal{V}$ are defined by

$$i_l^+ = \begin{cases} i_{l+1} & : 1 \leq l < L \\ * & : l = L \end{cases} \quad \text{and} \quad i_l^- = \begin{cases} i_{l-1} & : 1 < l \leq L \\ * & : l = 1 \end{cases}.$$

We assume that \mathcal{W} is another chain. The chains \mathcal{V} and \mathcal{W} induce a partial order on the product $\mathcal{U} = \mathcal{V} \times \mathcal{W}$ by

$$(i, j) \leq_{\mathcal{U}} (r, s) \Leftrightarrow i \leq_{\mathcal{V}} r \text{ and } j \leq_{\mathcal{W}} s$$

for all $(i, j), (r, s) \in \mathcal{U}$.

Definition 3.2. *Let $\mathcal{U} = \mathcal{V} \times \mathcal{W}$ be the product of chains \mathcal{V} and \mathcal{W} . The successor map on \mathcal{U} is a point-to-set map*

$$S_{\mathcal{U}} : \mathcal{U} \rightarrow 2^{\mathcal{U}}, \quad (i, j) \mapsto \{(i^+, j), (i, j^+), (i^+, j^+)\} \cap (\mathcal{V} \times \mathcal{W}),$$

where $2^{\mathcal{U}}$ denotes the set of all subsets of \mathcal{U} .

The successor map models the set of feasible warping steps for a given element $(i, j) \in \mathcal{U}$. Intersection of the successor map with $\mathcal{V} \times \mathcal{W}$ ensures that elements with $i^+ = *$ or $j^+ = *$ are excluded. The successor map sends (i, j) to the empty set if i and j are the last elements of the respective chains \mathcal{V} and \mathcal{W} . The next result shows that the successor map preserves the partial product order $\leq_{\mathcal{U}}$ as well as the linear orders $\leq_{\mathcal{V}}$ and $\leq_{\mathcal{W}}$.

Lemma 3.3. Let $\mathcal{U} = \mathcal{V} \times \mathcal{W}$ be the product of chains \mathcal{V} and \mathcal{W} . Suppose that $e = (i, j) \in \mathcal{U}$ is an element with $S_{\mathcal{U}}(e) \neq \emptyset$. Then the following order preserving properties hold:

1. $e \leq_{\mathcal{U}} e'$ for all $e' \in S_{\mathcal{U}}(e)$.
2. $i \leq_{\mathcal{V}} r$ and $j \leq_{\mathcal{W}} s$ for all $(r, s) \in S_{\mathcal{U}}(i, j)$.

Proof. Directly follows from the definitions of $\leq_{\mathcal{U}}$ and $S_{\mathcal{U}}$. \square

Lemma 3.4. Let $\mathcal{U} = \mathcal{V} \times \mathcal{W}$ be the product of chains \mathcal{V} and \mathcal{W} . Let $\mathcal{E} \subseteq \mathcal{U}$ be a subset consisting of L elements $e_1, \dots, e_L \in \mathcal{E}$ such that $e_{l+1} \in S_{\mathcal{U}}(e_l)$ for all $l \in [L-1]$. Then \mathcal{E} is a chain.

Proof. The successor map is order preserving with respect to the product order $\leq_{\mathcal{U}}$ according to Lemma 3.3. The assertion follows, because any order is transitive. \square

The chain $\mathcal{E} \subseteq \mathcal{U}$ in Lemma 3.4 is compatible with the successor map $S_{\mathcal{U}}$. We call such a chain a warping chain.

Definition 3.5. Let $\mathcal{U} = \mathcal{V} \times \mathcal{W}$ be the product of chains \mathcal{V} and \mathcal{W} . A warping chain in \mathcal{U} is a chain $\mathcal{E} = \{e_1, \dots, e_L\} \subseteq \mathcal{W}$ such that $e_{l+1} \in S_{\mathcal{U}}(e_l)$ for all $l \in [L-1]$.

The next result shows that warping chains preserve the order of the factor chains.

Proposition 3.6. Let \mathcal{V} and \mathcal{W} be chains. Let \mathcal{E} be a warping chain in $\mathcal{V} \times \mathcal{W}$. Then any pair of elements $(i, j), (r, s) \in \mathcal{E}$ satisfies

$$(i \leq_{\mathcal{V}} r \wedge j \leq_{\mathcal{W}} s) \vee (r \leq_{\mathcal{V}} i \wedge s \leq_{\mathcal{W}} j). \quad (3)$$

Proof. Suppose that $\mathcal{E} = \{e_1, \dots, e_L\}$. Then there are indices $p, q \in [L]$ such that $e_p = (i, j)$ and $e_q = (r, s)$. Without loss of generality, we assume that $p \leq q$. Let $u = q - p$. Repeatedly applying Lemma 3.3 yields

$$e_p \leq_{\mathcal{U}} \dots \leq_{\mathcal{U}} e_{p+u} = e_q,$$

where $\mathcal{U} = \mathcal{V} \times \mathcal{W}$. Since any order is transitive, we have $e_p \leq_{\mathcal{U}} e_q$. Then the assertion directly follows from the definition of the product order $\leq_{\mathcal{U}}$. \square

Equation (3) is the order-preserving property (or non-crossing property) of a warping chain. Note that the order preserving property does not hold for all subsets of $\mathcal{V} \times \mathcal{W}$.

3.2 Warping Graphs

This section introduces the notion of warping graph that models the alignment of two time series by a warping path and studies its local and global structure.

A graph is a pair $G = (\mathcal{V}, \mathcal{E})$ consisting of a finite set $\mathcal{V} \neq \emptyset$ of nodes and a set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ of edges. A node $i \in \mathcal{V}$ is incident with an edge $e \in \mathcal{E}$, if there is a node $j \in \mathcal{V}$ such that $e = (i, j)$ or $e = (j, i)$. Similarly, an edge $(i, j) \in \mathcal{E}$ is said to be incident to node i and to node j . The neighborhood of node $i \in \mathcal{V}$ is the subset of nodes defined by $\mathcal{N}(i) = \{j \in \mathcal{V} : (i, j) \in \mathcal{E} \text{ or } (j, i) \in \mathcal{E}\}$. The elements of $\mathcal{N}(i)$ are the neighbors of i . The degree $\deg(i) = |\mathcal{N}(i)|$ of node i in G is the number of neighbors of i .

A subgraph of graph $G = (\mathcal{V}, \mathcal{E})$ is a graph $G' = (\mathcal{V}', \mathcal{E}')$ such that $\mathcal{V}' \subseteq \mathcal{V}$ and $\mathcal{E}' \subseteq \mathcal{E}$. We write $G' \subseteq G$ to denote that G' is a subgraph of G . A graph G is connected, if for any two nodes $i, j \in \mathcal{V}$ there is a sequence $i = u_1, u_2, \dots, u_n = j$ of nodes in G such that $u_{k+1} \in \mathcal{N}(u_k)$ for all $k \in [n-1]$. A component C of graph G is a connected subgraph $C \subseteq G$ such that $C \subseteq C'$ implies $C = C'$ for every connected subgraph $C' \subseteq G$.

A graph $G = (\mathcal{U}, \mathcal{E})$ is bipartite, if \mathcal{U} can be partitioned into two disjoint and non-empty subsets \mathcal{V} and \mathcal{W} such that $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{W}$. We write $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ to denote a bipartite graph with node partitions \mathcal{V} and \mathcal{W} . Note that the order of the node partitions \mathcal{V} and \mathcal{W} in a bipartite graph $G = (\mathcal{U}, \mathcal{E})$ matters. A bipartite chain graph is a bipartite graph whose node partitions are chains.

Definition 3.7. A bipartite chain graph $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ with node partitions $\mathcal{V} = \{i_1, \dots, i_m\}$ and $\mathcal{W} = \{j_1, \dots, j_n\}$ is a warping graph of size $m \times n$ if



Figure 4: Example of a non-compact warping graph G (a) and its compactification H (b).

1. $(i_1, j_1), (i_m, j_m) \in \mathcal{E}$ (boundary condition)
2. \mathcal{E} is a warping chain in $\mathcal{V} \times \mathcal{W}$ (step condition),

The set of all warping graphs of size $m \times n$ is denoted by $\mathcal{G}_{m,n}$. If $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ is a warping graph, we briefly write S_G to denote the successor map $S_{\mathcal{V} \times \mathcal{W}}$ and \leq_G to denote the induced product order $\leq_{\mathcal{V} \times \mathcal{W}}$. The following result is a direct consequence of the boundary and step conditions:

Proposition 3.8. *Every node in a warping graph has a neighbor.*

We show that the neighborhood of a node of one partition of a warping graph is a contiguous chain of the other partition.

Proposition 3.9. *Let G be a warping graph with node partitions \mathcal{Z} and \mathcal{Z}' . Suppose that $i \in \mathcal{Z}$ is a node. Then the neighborhood $\mathcal{N}(i)$ of a node in $i \in \mathcal{Z}$ is a contiguous subchain of \mathcal{Z}' .*

Proof. Suppose that the warping graph is of the form $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$. Without loss of generality, we assume that $\mathcal{Z} = \mathcal{V}$ and $\mathcal{Z}' = \mathcal{W}$. Then we have $i \in \mathcal{V}$. The assertion trivially holds for $|\mathcal{N}(i)| = 1$. Suppose that $|\mathcal{N}(i)| > 1$. We assume that $\mathcal{N}(i)$ is not contiguous. Then there are elements $j', j'' \in \mathcal{N}(i)$ and $j \in \mathcal{W} \setminus \mathcal{N}(i)$ such that $j' \leq_G j \leq_G j''$. From Prop. 3.8 follows that there is a node $i' \in \mathcal{V} \setminus \{i\}$ such that $(i', j) \in \mathcal{E}$.

Two cases can occur: (1) $i' <_{\mathcal{V}} i$ and (2) $i <_{\mathcal{V}} i'$. It is sufficient to consider the first case $i' <_{\mathcal{V}} i$. The proof of the second case is analogue. By construction, there are edges (i', j) and (i, j') such that $i' <_{\mathcal{V}} i$ and $j' <_{\mathcal{W}} j$. These relationships violate the order preserving property of a warping chain given in Eq. (3) of Prop. 3.6. Hence, $\mathcal{N}(i)$ is contiguous. \square

We introduce compact warping graphs that represent warping paths of minimal length.

Definition 3.10. *A warping graph $G \in \mathcal{G}_{m,n}$ is compact if there is no warping graph $G' \in \mathcal{G}_{m,n}$ such that G' is a proper subgraph of G .*

A warping graph is compact if no edge can be deleted without violating the boundary or step conditions. Figure 4 shows an example of a non-compact warping graph and its compactification.

Proposition 3.11. *Let G be a warping graph with edge set $\mathcal{E} = \{e_1, \dots, e_L\}$. Then the following statements are equivalent:*

1. G is compact.
2. Let $2 \leq k < L$. Then $e_{l+k} \notin S_G(e_l)$ for all $l \in [L - k]$.

Proof. We first prove the following Lemma:

Lemma. *Let \mathcal{V} and \mathcal{W} be two chains, let $\mathcal{E} = \{e_1, \dots, e_L\} \subseteq \mathcal{V} \times \mathcal{W}$ be a warping chain, and let $3 \leq k < L$. Then $e_{l+k} \notin S_{\mathcal{V} \times \mathcal{W}}(e_l)$ for every $l \in [L - k]$.*

Proof. Let $\mathcal{C} = \{i_1, \dots, i_m\}$ be a chain and let $i_k, i_l \in \mathcal{C}$ be two elements of \mathcal{C} . Then we define the distance

$$\Delta_{\mathcal{C}}(i_k, i_l) = |l - k| + 1.$$

Suppose that $e_l = (i, j)$ for some $l \in [L - 1]$. Let $e' = (r, s) \in S_{\mathcal{V} \times \mathcal{W}}(s_l)$ be an arbitrary successor of e_l . Then by definition of the successor map, we have $\Delta_{\mathcal{V}}(i, r), \Delta_{\mathcal{W}}(j, s) \in \{0, 1\}$. Let $l \in [L - k]$. Suppose that $e_l = (i, j)$ and $e_{l+k} = (r, s)$. Then by induction, we have

$$\Delta_{\mathcal{V}}(i, r) + \Delta_{\mathcal{W}}(j, s) \geq k.$$

From $k \geq 3$ follows $\Delta_{\mathcal{V}}(i, r) \geq 2$ or $\Delta_{\mathcal{W}}(j, s) \geq 2$. Hence, $\Delta_{\mathcal{V}}(i, r) \notin \{0, 1\}$ or $\Delta_{\mathcal{W}}(j, s) \notin \{0, 1\}$. This shows the assertion $e_{l+k} \notin S_{\mathcal{V} \times \mathcal{W}}(e_l)$. \square

From the above Lemma follows that the case $k > 2$ is impossible for a warping chain. Therefore it is sufficient to consider the case $k = 2$.

Let G be compact. We assume that there is an $l \in [L - 2]$ such that $e_{l+2} \in S_G(e_l)$. By construction, the edge e_{l+1} is an inner element of the chain \mathcal{E} . From $e_{l+2} \in S_G(e_l)$ follows that removing e_{l+1} neither violates the boundary conditions nor the step condition. This contradicts compactness of G and shows that a compact warping graph G implies the second statement.

Next, we show the opposite direction. Suppose that $e_{l+2} \notin S_G(e_l)$ for all $l \in [L - 2]$. We assume that G is not compact. Then there is an edge $e_k \in \mathcal{E}$ that can be removed without violating the boundary and step conditions. Not violating the boundary condition implies that $1 < k < L$. Hence, e_{k-1} and e_{k+1} are edges in \mathcal{E} . We set $l = k - 1$. Then we obtain the contradiction that $e_l \in S_G(e_{l+2})$. Hence, G is compact. \square

Suppose that $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ is a warping graph. By $\mathcal{V} \sqcup \mathcal{W}$ we denote the disjoint union of the node partitions. If $i \in \mathcal{V} \sqcup \mathcal{W}$ is a node of one partition, then its neighborhood $\mathcal{N}(i)$ is a subset of the other node partition. Hence, $\mathcal{N}(i)$ is a chain and has boundary and eventually inner nodes. Let $\mathcal{N}^\circ(i)$ denote the possibly empty subset of inner nodes of chain $\mathcal{N}(i)$. We show that inner nodes of $\mathcal{N}(i)$ always have degree one.

Lemma 3.12. *Let $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ be a warping graph. Suppose that $i \in \mathcal{V} \sqcup \mathcal{W}$ is a node with neighborhood $\mathcal{N}(i)$. Then $\deg(j) = 1$ for all $j \in \mathcal{N}^\circ(i)$.*

Proof. Without loss of generality, we assume that $i \in \mathcal{V}$. Then $\mathcal{N}(i) \subseteq \mathcal{W}$ is a chain. The assertion holds for $|\mathcal{N}(i)| \leq 2$, because in this case $\mathcal{N}(i)$ has no inner node. Suppose that $|\mathcal{N}(i)| > 2$. Let $j \in \mathcal{N}(i)$ be an inner node. We assume that $\deg(j) > 1$. Then there is a node $i' \in \mathcal{V} \setminus \{i\}$ such that $(i', j) \in \mathcal{E}$. Since \mathcal{V} is a chain, we find that either $i' <_{\mathcal{V}} i$ or $i <_{\mathcal{V}} i'$.

We only consider the first case $i' <_{\mathcal{V}} i$. The proof of the second case is analogue. Observe that j is an inner node of $\mathcal{N}(i)$ and $\mathcal{N}(i)$ is contiguous by Prop. 3.9. Then $\mathcal{N}(i)$ contains the predecessor $j' = j^-$ of node j . By construction $e = (i, j')$ and $e' = (i', j)$ are edges of \mathcal{E} such that $i' <_{\mathcal{V}} i$ and $j' <_{\mathcal{W}} j$. Thus, the edges e and e' violate the order preserving property of a warping chain given in Eq. (3) of Prop. 3.6. This contradicts our assumption that \mathcal{E} is a warping chain. Hence, we have $\deg(j) = 1$. Since the inner node j was chosen arbitrarily, the assertion follows. \square

Lemma 3.13. *Let $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ be a warping graph and let $i \in \mathcal{V} \sqcup \mathcal{W}$ be a node with neighborhood $\mathcal{N}(i)$. Suppose that $|\mathcal{N}(i)| \geq 2$ and $j \in \mathcal{N}(i)$ is a boundary node with $\deg(j) \geq 2$. Then the following properties hold:*

1. *If j is the first node in $\mathcal{N}(i)$, then $i^- \in \mathcal{V}$ exists and $(i^-, j) \in \mathcal{E}$.*
2. *If j is the last node in $\mathcal{N}(i)$, then $i^+ \in \mathcal{V}$ exists and $(i^+, j) \in \mathcal{E}$.*

Proof. We show the second assertion. The proof of the first assertion is analogue. Since $|\mathcal{N}(i)| \geq 2$ and j is the last node of $\mathcal{N}(i)$, we find that $j' = j^- \in \mathcal{N}(i)$ and therefore $(i, j') \in \mathcal{E}$ exists. From $\deg(j) \geq 2$ follows that there is a node $i' \in \mathcal{V}$ such that $(i', j) \in \mathcal{E}$. We assume that $i' \in \mathcal{N}(j)$ satisfies $i' <_{\mathcal{V}} i$. This implies that $e = (i, j')$ and $e' = (i', j)$ are two edges of \mathcal{E} such that $i' <_{\mathcal{V}} i$ and $j' <_{\mathcal{W}} j$. Thus, the edges e and e' violate the order preserving property of a warping chain given in Eq. (3) of Prop. 3.6. This contradicts our assumption that $i' <_{\mathcal{V}} i$. Therefore, we have $i <_{\mathcal{V}} i'$. This in turn shows that $i^+ \in \mathcal{V}$ exists. From $i, i' \in \mathcal{N}(j)$ and $i <_{\mathcal{V}} i'$ follows $i^+ \in \mathcal{N}(j)$, because $\mathcal{N}(j)$ is a contiguous subchain of \mathcal{V} by Prop. 3.9. This shows $(i^+, j) \in \mathcal{E}$ and completes the proof. \square

A bipartite graph $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ is complete if $\mathcal{E} = \mathcal{V} \times \mathcal{W}$. Let $r \in \mathbb{N}$. A complete bipartite graph $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ is a star graph of the form $K_{1,r}$, if $|\mathcal{V}| = 1$ and $|\mathcal{W}| = r$. Similarly, G is a star graph of the form $K_{r,1}$, if $|\mathcal{V}| = r$ and $|\mathcal{W}| = 1$. By definition, a star graph has at least two nodes. A star forest is a graph whose components are star graphs.

Proposition 3.14. *A compact warping graph is a star forest.*

Proof. Let $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ be a compact warping graph. Suppose that $C = (\mathcal{V}', \mathcal{W}', \mathcal{E}')$ is a component of G . From Prop. 3.8 follows that C has at least two nodes connected by an edge.

We assume that C is not a star. Then C has two nodes $i, j \in \mathcal{V}' \sqcup \mathcal{W}'$ with degree larger than one. Without loss of generality, we assume that $i \in \mathcal{V}'$. Then $\mathcal{N}(i) \subseteq \mathcal{W}'$ has at least two elements. Suppose that all nodes from $\mathcal{N}(i)$ have degree one. Since component C is bipartite, we find that C is isomorphic to the star $K_{1,r}$, where $r = \deg(i) > 1$. This contradicts our assumption that C is not a star. Hence, there is a node $j \in \mathcal{N}(i) \subseteq \mathcal{W}'$ with $\deg(j) > 1$.

From Lemma 3.12 follows that node j is a boundary node of $\mathcal{N}(i)$. We show the assertion for the case that j is the last node in $\mathcal{N}(i)$. The proof for the case that j is the first node in $\mathcal{N}(i)$ is analogue. Since j is the last node in $\mathcal{N}(i)$ and $|\mathcal{N}(i)| \geq 2$, we have $j^- \in \mathcal{N}(i)$ and therefore $(i, j^-) \in \mathcal{E}$. Applying Lemma 3.13 yields that $i^+ \in \mathcal{V}$ exists and $(i^+, j) \in \mathcal{E}$.

By construction, we have $(i, j^-), (i, j), (i^+, j) \in \mathcal{E}$. This shows that (i, j) is not a boundary edge in \mathcal{E} . Since $(i^+, j) \in S_{\mathcal{V} \times \mathcal{W}}(i, j^-)$, we can remove (i, j) without violating the step condition. Then the subgraph $G' = (\mathcal{V}, \mathcal{W}, \mathcal{E} \setminus \{(i, j)\})$ of G is a warping graph. This contradicts our assumption that G is compact. Hence, C is a star. \square

An immediate consequence of the proof of Prop. 3.14 is the following corollary.

Corollary 3.15. *Let $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ be a compact warping graph in $\mathcal{G}_{m,n}$. Then every component of G is a star with at least one node in \mathcal{V} and one node in \mathcal{W} .*

Proposition 3.16. *Let $G \in \mathcal{G}_{m,n}$ be a compact warping graph with $m > n$. Then we have:*

1. *G has at most $n - 1$ components of the form $K_{1,1}$.*
2. *G has a component of the form $K_{r,1}$ with $r > 1$.*

Proof. Let $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ with $|\mathcal{V}| = m$ and $|\mathcal{W}| = n$. To show the first assertion, we assume that G has $n' > n - 1$ components of the form $K_{1,1}$. This is only possible for $n' = n$, because every component of G has at least one node in \mathcal{W} by Corollary 3.15.

Let C_1, \dots, C_n be n components of G of the form $K_{1,1}$. Suppose that $\mathcal{C}_k = (\{i_k\}, \{j_k\}, \{(i_k, j_k)\})$ for all $k \in [n]$. The union of the first node partitions over the n components C_k gives $\mathcal{V}' = \{i_1, \dots, i_n\}$. From $m > n$ follows $\mathcal{V}' \subsetneq \mathcal{V}$. Then there is a node $i \in \mathcal{V} \setminus \mathcal{V}'$. From Prop. 3.14 follows that there is a component C of G that is a star of the form $K_{r,s}$ that contains node i . Since $s \geq 1$ by definition of a star, component C has a node $j \in \mathcal{W}$. Then there is a $k \in [n]$ such that $j = j_k$ is a node in component C_k . Since $i \neq i_k$ by construction, the graph $H = (\{i, i_k\}, \{j_k\}, \{(i, j_k), (i_k, j_k)\})$ is a connected subgraph of G that includes component C_k as a proper subgraph. This contradicts our assumption that C_k is a maximal connected subgraph of G . Consequently, G cannot have more than $n - 1$ components of the form $K_{1,1}$.

Next, we show the second assertion. Suppose that C_1, \dots, C_q are all components of G that are of the form $K_{1,1}$. Let $\mathcal{V}' = \{i_1, \dots, i_q\} \subseteq \mathcal{V}$ and $\mathcal{W}' = \{j_1, \dots, j_q\} \subseteq \mathcal{W}$ be the subsets covered by the q components C_k . From the first part of this proof follows that $q < n$ and by assumption, we have $n < m$. Then $\mathcal{V}'' = \mathcal{V} \setminus \mathcal{V}'$ and $\mathcal{W}'' = \mathcal{W} \setminus \mathcal{W}'$ are non-empty. By $m'' = |\mathcal{V}''| = m - q$ and $n'' = |\mathcal{W}''| = n - q$ we denote the respective number of nodes not contained in any of the q components C_k . From $q < n < m$ follows that $1 \leq n'' < m''$. The pigeonhole principle states that there is at least one node $j \in \mathcal{W}''$ that is connected to at least two nodes $i, i' \in \mathcal{V}''$. Let C be the component of G containing the three nodes i, i' , and j . From Prop. 3.14 follows that C is a star of the form $K_{r,s}$. We find that $r \geq 2$, because C contains at least the two nodes i and i' from $\mathcal{V}'' \subset \mathcal{V}$. Then $s = 1$, because C is a star. This shows the second assertion. \square

Definition 3.17. *A node i of a compact warping graph G is redundant if $\deg(j) \geq 2$ for all $j \in \mathcal{N}(i)$.*

Let i be a node in G . Then $G - \{i\}$ is the subgraph of G obtained by deleting node i and its incident edges. The next result shows that deleting a redundant node of a compact warping graph is again a compact warping graph.

Proposition 3.18. *Let i be a redundant node of a compact warping graph G . Then $G - \{i\}$ is a compact warping graph.*

Proof. Without loss of generality, we assume that $i \in \mathcal{V}$. Let $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ with warping chain $\mathcal{E} = \{e_1, \dots, e_L\}$. Suppose that $G' = G - \{i\} = (\mathcal{V}', \mathcal{W}, \mathcal{E}')$, where $\mathcal{V}' = \mathcal{V} \setminus \{i\}$ and \mathcal{E}' is the chain obtained from \mathcal{E} by removing all edges incident to node i .

We first show that $\mathcal{N}(i) \subseteq \mathcal{W}$ consists of a singleton. Let C be the component of G that contains node i . Then C also includes the neighborhood $\mathcal{N}(i)$. Since G is compact, we can apply Prop. 3.14 and find that component C is a star of the form $K_{r,1}$ or $K_{1,r}$, where $r \geq 1$. Observe that $r = \deg(j) \geq 2$ for every neighbor $j \in \mathcal{N}(i)$, because node i is redundant. Hence, C is a star of the form $K_{r,1}$ and therefore $\mathcal{N}(i) = \{j\}$.

From $\deg(j) \geq 2$ follows that there is a node $i' \in \mathcal{V} \setminus \{i\}$ such that $(i', j) \in \mathcal{E}$. We distinguish between two cases: (1) $i' <_{\mathcal{V}} i$ and (2) $i <_{\mathcal{V}} i'$. We only consider the first case $i' <_{\mathcal{V}} i$. The proof of the second case is analogue. From Prop. 3.9 follows that $\mathcal{N}(j)$ is a contiguous subchain of \mathcal{W} with $i', i \in \mathcal{N}(j)$. Then $i^- \in \mathcal{N}(j)$. We distinguish between two cases: (1) $e_L = (i, j)$, and (2) $e_l = (i, j)$ for some $1 < l < L$. Note that the case $e_1 = (i, j)$ cannot occur due to existence of $i^- \in \mathcal{V}$.

Case 1: From $e_L = (i, j)$ follows that the edge set of G' is of the form $\mathcal{E}' = \{e_2, \dots, e_{L-1}\}$. In addition, i is the last node in \mathcal{V} and therefore i^- is the last node in \mathcal{V}' . Furthermore, we find that $j \in \mathcal{W}$ is the last node in \mathcal{W} . We show that \mathcal{E}' is a warping chain. The first boundary condition is satisfied by $e_1 \in \mathcal{E}'$. From $i^- \in \mathcal{N}(j)$ follows that $(i^-, j) \in \mathcal{E}$ is an edge that satisfies the second boundary condition connecting the last nodes of \mathcal{V}' and \mathcal{W} . Finally, from $e_{l+1} \in S_G(e_l)$ for all $l \in [L-1]$ follows that the step condition remains valid in G' . Therefore, the edge set \mathcal{E}' is a warping chain of length $L' = L - 1$. It remains to show that G' is compact. For this, we assume that G' is not compact. Then from Prop. 3.11 follows that there is an index $l \in [L' - 2]$ such that $e_{l+2} \in S_{G'}(e_l)$. This implies that $e_{l+2} \in S_G(e_l)$ contradicting the assumption that G is compact. Hence, G' is a compact warping graph.

Case 2: There is an index $1 < l < L$ such that $e_l = (i, j)$. Hence, \mathcal{E} has at least three edges and the edge set of G' is of the form $\mathcal{E}' = \{e_1, \dots, e_{l-1}, e_{l+1}, \dots, e_L\}$. To show that \mathcal{E}' is a warping chain, we assume that $e_{l-1} = (i', j')$ and $e_{l+1} = (i'', j'')$. From $\mathcal{N}(i) = \{j\}$ and the step condition follows that $i' = i^-$ and $i'' = i^+$. This shows that i is neither the first nor last node in \mathcal{V} . Hence, e_1 and e_L satisfy the boundary conditions in \mathcal{E}' .

We show that \mathcal{E}' satisfies the step condition. Since \mathcal{E} is a warping chain, we have $e_{k+1} = S_G(e_k)$ for all $k \in [L-1]$. Since $S_{G'} = S_G$ on \mathcal{E}' , it is sufficient to show that $e_{l+1} \in S_{G'}(e_l)$. According to the previous parts of the proof, we have $(i^-, j) \in \mathcal{E}$. The step condition together with $i' = i^-$ imply that $e_{l-1} = (i', j') = (i^-, j') = (i^-, j)$ and therefore $j' = j$. Again, from the step condition follows that either $j'' = j$ or $j'' = j^+$. Observe that $i'^+ = i''$ in \mathcal{V}' . Then we have $(i'', j) \in S_{G'}(i', j)$ and $(i'', j^+) \in S_{G'}(i', j)$. This shows that \mathcal{E}' satisfies the step condition in either of both cases $j'' = j$ or $j'' = j^+$.

It remains to show that G' is compact. For this, we assume that G' is not compact. Suppose that $\mathcal{I} = [L] \setminus \{l\}$ is the index set of \mathcal{E}' . Then from Prop. 3.11 follows that there is an index $k \in \mathcal{I} \setminus \{L-1, L\}$ such that $e_{k+2} \in S_{G'}(e_k)$. This implies that $e_{k+2} \in S_G(e_k)$ contradicting the assumption that G is compact. Hence, G' is a compact warping graph. \square

3.3 Glued Warping Graphs

This section glues warping graphs to model Fréchet function. Then the graph-theoretic foundation of the Reduction Theorem is presented.

A graph $G = (\mathcal{U}, \mathcal{E})$ is a centered N -partite graph if the set \mathcal{U} can be partitioned into $N + 1$ disjoint and non-empty subsets $\mathcal{V}, \mathcal{W}_1, \dots, \mathcal{W}_N$ such that

$$\mathcal{E} \subseteq \bigcup_{k=1}^N \mathcal{V} \times \mathcal{W}_k.$$

Definition 3.19. Let G_1, \dots, G_N be compact warping graphs of the form $G_k = (\mathcal{V}, \mathcal{W}_k, \mathcal{E}_k)$ for all $k \in [N]$. The glued graph with splice \mathcal{V} and particles G_1, \dots, G_N is a centered N -partite graph $G = (\mathcal{V}, \mathcal{W}_1, \dots, \mathcal{W}_N, \mathcal{E})$ with edge set $\mathcal{E} = \mathcal{E}_1 \sqcup \dots \sqcup \mathcal{E}_N$.

Note that a particle of a glued graph is always a compact warping graph. The definition of a glued graph assumes that all N particles G_1, \dots, G_N share a common node partition \mathcal{V} and that any two particles G_k and G_l have disjoint node partitions \mathcal{W}_k and \mathcal{W}_l , respectively. Then the glued graph with splice \mathcal{V} is obtained by taking the disjoint union of the particles G_1, \dots, G_N , but by identifying the nodes from \mathcal{V} . A special case of a glued graph is any compact warping graph $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ with the first partition \mathcal{V} as its splice.

Definition 3.20. Let G be a glued graph with splice \mathcal{V} and particles G_1, \dots, G_N . Node $i \in \mathcal{V}$ is *redundant* in G , if it is redundant in G_k for every $k \in [N]$.

Proposition 3.21. Let G be a glued graph with splice \mathcal{V} and particles G_1, \dots, G_N . Suppose that $i \in \mathcal{V}$ is redundant. Then $G - \{i\}$ is a glued graph with splice $\mathcal{V} \setminus \{i\}$ and particles $G_1 - \{i\}, \dots, G_N - \{i\}$.

Proof. Let $k \in [N]$. Then particle $G_k = (\mathcal{V}, \mathcal{W}_k, \mathcal{E}_k)$ is a compact warping graph by definition of a glued graph. Since splice node $i \in \mathcal{V}$ is redundant in G , it is redundant in G_k . Then from Prop. 3.18 follows that $G'_k = G_k - \{i\} = (\mathcal{V}', \mathcal{W}_k, \mathcal{E}')$ is a compact warping graph with $\mathcal{V}' = \mathcal{V} \setminus \{i\}$ and $\mathcal{E}'_k = \mathcal{E}_k \setminus \mathcal{E}_k(i)$, where $\mathcal{E}_k(i) \subseteq \mathcal{E}_k$ is the subset of edges in G_k incident to node i .

The graph $G' = G - \{i\} = (\mathcal{V}', \mathcal{W}_1, \dots, \mathcal{W}_N, \mathcal{E}')$ has an edge set of the form $\mathcal{E}' = \mathcal{E} \setminus \mathcal{E}(i)$, where $\mathcal{E}(i) \subseteq \mathcal{E}$ is the subset of edges in G incident to node i . Since $\mathcal{E} = \mathcal{E}_1 \sqcup \dots \sqcup \mathcal{E}_N$, we have

$$\mathcal{E}(i) = \mathcal{E}_1(i) \sqcup \dots \sqcup \mathcal{E}_N(i) = \mathcal{E}'_1 \sqcup \dots \sqcup \mathcal{E}'_N.$$

This shows that $G - \{i\}$ is a glued graph of particles G'_1, \dots, G'_N along splice \mathcal{V}' . \square

Suppose that G is a glued graph with splice \mathcal{V} and particles G_1, \dots, G_N . A particle is said to be *trivial* if it is a star of the form $K_{m,1}$. By $\mathcal{I}_G \subseteq [N]$ we denote the subset of all indices $k \in [N]$ for which particle G_k is non-trivial. We call \mathcal{I}_G the *core index set* (core) of G .

Definition 3.22. Let G be a glued graph with splice \mathcal{V} and N particles $G_k = (\mathcal{V}, \mathcal{W}_k, \mathcal{E}_k)$ for all $k \in [N]$. Suppose that \mathcal{I}_G is the core index set of G . Then

$$\rho(G) = \begin{cases} \sum_{k \in \mathcal{I}_G} |\mathcal{W}_k| - 2(|\mathcal{I}_G| - 1) & : \mathcal{I}_G \neq \emptyset \\ 1 & : \mathcal{I}_G = \emptyset \end{cases}$$

is the reduction bound of G .

Now we present the graph-theoretic foundation of the Reduction Theorem.

Theorem 3.23. Let G be a glued graph with splice \mathcal{V} such that $\rho(G) < |\mathcal{V}|$. Then \mathcal{V} has a redundant node.

Proof. Suppose that G_1, \dots, G_N are the particles of G with $G_k = (\mathcal{V}, \mathcal{W}_k, \mathcal{E}_k)$ for all $k \in [N]$. Let $m = |\mathcal{V}|$, $n_k = |\mathcal{W}_k|$, and $\mathcal{N}_k(i) = \mathcal{N}(i) \cap \mathcal{W}_k$ for every $k \in [N]$.

We first consider the special case that $\mathcal{I}_G = \emptyset$. Then all N particles are trivial, that is $n_k = 1$ for every $k \in [N]$. The reduction bound is of the form $\rho(G) = 1$. By assumption, we have $m > \rho(G)$. Hence, every particle G_k is a star of the form $K_{m,1}$, where $m > 1$. This shows that every splice node is redundant.

Next, we assume that $\mathcal{I}_G \neq \emptyset$. We set $N' = |\mathcal{I}_G|$. Obviously, we have $N' \geq 1$. We say, G_k supports node $i \in \mathcal{V}$, if there is a node $j \in \mathcal{N}_k(i) \subseteq \mathcal{W}_k$ with $\deg(j) = 1$. It is sufficient to show that \mathcal{V} has a node not supported by any of the N particles G_1, \dots, G_N . The proof proceeds in four steps.

1. We show that $n_k < m$ for any $k \in \mathcal{I}_G$. Suppose that $\mathcal{J} = \mathcal{I}_G \setminus \{k\}$. Then from $\mathcal{I}_G \neq \emptyset$ follows

$$\rho(G) = \sum_{l \in \mathcal{I}_G} n_l - 2(N' - 1) = n_k + \sum_{l \in \mathcal{J}} n_l - 2(N' - 1).$$

From $l \in \mathcal{I}_G$ follows $n_l \geq 2$. This together with $|\mathcal{J}| = N' - 1$ yields

$$\rho(G) \geq n_k + \sum_{l \in \mathcal{J}} 2 - 2(N' - 1) = n_k + 2|\mathcal{J}| - 2(N' - 1) = n_k.$$

Then from $m > \rho(G)$ follows $m > n_k \geq 2$.

2. We bound the number of splice nodes that can be supported by any non-trivial particle. For any $k \in \mathcal{I}_G$ let $\mathcal{W}'_k \subseteq \mathcal{W}_k$ be the subset of nodes in G_k that support a splice node. We define a map $\phi_k : \mathcal{W}'_k \rightarrow \mathcal{V}$ such that $(\phi_k(j), j) \in \mathcal{E}_k$. Such a map exists due to the boundary and step conditions of warping graph G_k . Moreover, the map ϕ_k is uniquely determined, because $\deg(j) = 1$ for any node $j \in \mathcal{W}'_k$. This shows that $\mathcal{V}_k = \phi_k(\mathcal{W}'_k)$ is the set of splice nodes supported by G_k . Since ϕ_k is bijective, we have $|\mathcal{V}_k| = |\mathcal{W}'_k|$.

From step 1 of this proof follows that G_k is a compact warping graph with $n_k < m$. Therefore, we can apply Prop. 3.16 and obtain that G_k has at most $n_k - 1$ components of the form $K_{1,1}$ and at least once component of the form $K_{r,1}$ with $r > 1$. This shows that $|\mathcal{W}'_k| = |\mathcal{V}_k| \leq n_k - 1$.

3. We show that there is a splice node not supported by any non-trivial particle. For this, we define the set

$$\mathcal{U} = \bigcup_{k \in \mathcal{I}_G} \mathcal{V}_k$$

of all splice nodes that are supported by at least one non-trivial particle of G . Then it is sufficient to show that $m > |\mathcal{U}|$. We consider three cases: (1) $N' = 1$, (2) $N' = 2$, and (3) $N' > 2$.

Case 1: $N' = 1$. Suppose that $\mathcal{I}_G = \{u\}$. Since $\mathcal{I}_G \neq \emptyset$, the reduction bound is of the form

$$\rho(G) = n_u - 2(N' - 1) = n_u \geq 2.$$

According to step 2, we have $n_u - 1 \geq |\mathcal{V}_u|$. By using $\mathcal{U} = \mathcal{V}_u$ we find that

$$m > \rho(G) > |\mathcal{V}_u| = |\mathcal{U}|.$$

Case 2: $N' = 2$. Suppose that $\mathcal{I}_G = \{u, v\}$. Since $\mathcal{I}_G \neq \emptyset$, the reduction bound takes the form

$$\rho(G) = n_u + n_v - 2(N' - 1) = n_u + n_v - 2 = (n_u - 1) + (n_v - 1).$$

According to step 2, we have $n_u - 1 \geq |\mathcal{V}_u|$ and $n_v - 1 \geq |\mathcal{V}_v|$. By using $\mathcal{U} = \mathcal{V}_u \cup \mathcal{V}_v$ we obtain

$$m > \rho(G) \geq |\mathcal{V}_u| + |\mathcal{V}_v| \geq |\mathcal{U}|.$$

Case 3: $N' > 2$. Suppose that the slice \mathcal{V} is a chain of the form $\mathcal{V} = \{i_1, \dots, i_m\}$ with boundary nodes $\text{bd}(\mathcal{V}) = \{i_1, i_m\}$. We assume that $|\mathcal{U}| = m$. Then there are (not necessarily distinct) indices $u, v \in \mathcal{I}_G$ such that $i_1 \in \mathcal{V}_u$ and $i_m \in \mathcal{V}_v$. From the boundary and step conditions follows that the boundary nodes of any W_k ($k \in \mathcal{I}_G$) can only support the respective boundary nodes of \mathcal{V} and not any other splice node. Then the first node of \mathcal{W}_u only supports $i_1 \in \mathcal{V}$ and the last node of \mathcal{W}_v only supports $i_m \in \mathcal{V}$.

Let $\mathcal{J} = \{u, v\}$, let $\mathcal{J}' = \mathcal{I}_G \setminus \mathcal{J}$, and let $\mathcal{V}'_k = \mathcal{V}_k \setminus (\mathcal{V}_u \cup \mathcal{V}_v)$ for all $k \in \mathcal{J}'$. The set \mathcal{V}'_k consists of all splice nodes supported by G_k but not by G_u and G_v . Hence, the boundary nodes of \mathcal{V} are not contained in \mathcal{V}'_k . Then both boundary nodes of \mathcal{W}_k do not support any node in \mathcal{V}'_k . This implies $|\mathcal{V}'_k| \leq n_k - 2$. From the cardinality of the set union follows

$$|\mathcal{U}| \leq \sum_{l \in \mathcal{J}} (n_l - 1) + \sum_{k \in \mathcal{J}'} (n_k - 2) = |\mathcal{J}| + \sum_{l \in \mathcal{J}} (n_l - 2) + \sum_{k \in \mathcal{J}'} (n_k - 2) = \sum_{k \in \mathcal{I}_G} n_k - 2N' + |\mathcal{J}|.$$

Since $|\mathcal{J}| \leq 2$, we obtain

$$|\mathcal{U}| \leq \sum_{k \in \mathcal{I}_G} n_k - 2(N' - 2) = \rho(G) < m.$$

This contradicts the assumption $|\mathcal{U}| = m$ and shows that $|\mathcal{U}| < m$ holds.

All three cases show that $|\mathcal{U}| < m$. Hence, G has a splice node not supported by any of the non-trivial particles.

4. We show that G has a splice node not supported by any of the trivial and non-trivial particles. The non-trivial part follows from step 3. Therefore, it is sufficient to consider trivial particles only. Since $\mathcal{I}_G \neq \emptyset$ by assumption, there is a $k \in \mathcal{I}_G$. From $m > n_k$ and $n_k \geq 2$ follows $m > 2$. This implies that the trivial particles of G do not support any of the splice nodes in \mathcal{V} . This completes the proof. \square

3.4 Labeled Warping Graphs

This section labels the nodes of warping graphs with the attributes of the corresponding time series to be aligned. Then we introduce the weight of a labeled warping graph for modeling the cost of aligning two time series along a warping path.

We assume that \mathcal{A} is an attribute set and $d : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}$ is a non-negative distance function on \mathcal{A} .

Definition 3.24. A labeled warping graph $H = (G, \lambda)$ consists of a warping graph $G = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ and a labeling function $\lambda : \mathcal{V} \sqcup \mathcal{W} \rightarrow \mathcal{A}$.

The labeling function λ assigns an attribute $\lambda(i) \in \mathcal{A}$ to any node $i \in \mathcal{V} \sqcup \mathcal{W}$. Thus, the nodes correspond to time points and the attributes to the elements at every time point.

The set of all labeled warping graphs of order $m \times n$ with label function λ is denoted by $\mathcal{G}_{m,n}^\lambda$. Since the set $\mathcal{G}_{m,n}^\lambda$ fixes both node partitions and the label function, the graphs in $\mathcal{G}_{m,n}^\lambda$ differ only in their edge sets. Thus, $\mathcal{G}_{m,n}^\lambda$ describes the set of all possible warping paths that align time series $x = (x_1, \dots, x_m)$ and $y = (y_1, \dots, y_n)$ whose elements $x_i = \lambda(i)$ and $y_j = \lambda(j)$ are specified by the labeling function λ .

Definition 3.25. Let $H = (G, \lambda)$ be a labeled warping graph with edge set \mathcal{E} . The weight of H is defined by

$$\omega(H) = \sum_{(i,j) \in \mathcal{E}} d(\lambda(i), \lambda(j))$$

The weight of a labeled warping graph corresponds to the cost of aligning two time series along a warping path. A DTW-graph is a labeled warping graph with minimal weight.

Definition 3.26. A graph $H \in \mathcal{G}_{m,n}^\lambda$ is a DTW-graph, if

$$\omega(H) = \min \{ \omega(H') : H' \in \mathcal{G}_{m,n}^\lambda \}.$$

The weight of a DTW-graph is the DTW-distance between the time series represented by the labeled node partitions.

3.5 Proofs of Results from Section 2

Proof of Theorem 2.7

Theorem (Reduction Theorem). Let F be the Fréchet function of a sample $\mathcal{X} \in \mathcal{T}^N$. Then for every time series $x \in \mathcal{T}$ of length $\ell(x) > \rho(\mathcal{X})$ there is a time series $x' \in \mathcal{T}$ of length $\ell(x') = \ell(x) - 1$ such that $F(x') \leq F(x)$.

Proof. Let $\mathcal{X} = (x^{(1)}, \dots, x^{(k)})$, $m = \ell(x)$, and $n_k = \ell(x^{(k)})$ for all $k \in [N]$. By assumption, we have $m > \rho(\mathcal{X})$.

For every $k \in [N]$ there is an optimal warping path $p^{(k)} \in \mathcal{P}_{m, n_k}$ aligning x and $x^{(k)}$. Let $H_k = (G_k, \lambda_k)$ be the DTW-graph representing $p^{(k)}$. Then $\omega(H_k) = \delta(x, x^{(k)})$ and $G_k = (\mathcal{V}, \mathcal{W}_k, \mathcal{E}_k) \in \mathcal{G}_{m, n_k}$ is a warping graph with $m = |\mathcal{V}|$ and $n_k = |\mathcal{W}_k|$. Then we have

$$F(x) = \frac{1}{N} \sum_{k=1}^N h(\omega(H_k)),$$

where h is the loss function.

Suppose that G_k is non-compact and $G'_k \subseteq G_k$ is compact. Since H_k is a DTW-graph, we have $\omega(H'_k) = \omega(H_k)$, where $H'_k = (G'_k, \lambda'_k)$. Hence, without loss of generality we can assume that G_k is compact for all $k \in [N]$. Let G be the glued graph with splice \mathcal{V} and particles G_1, \dots, G_N . Since $m > \rho(G)$, we can apply Prop. 3.23 and obtain that G has a redundant splice node $i \in \mathcal{V}$. Applying Prop. 3.21 yields that $G' = G - \{i\}$ is a glued graph with splice $\mathcal{V}' = \mathcal{V} \setminus \{i\}$ and particles G'_1, \dots, G'_N . The particles G'_k are of the form $G'_k = G_k - \{i\} = (\mathcal{V}', \mathcal{W}_k, \mathcal{E}'_k)$, where the edge set \mathcal{E}'_k is obtained from \mathcal{E}_k by removing all edges incident to splice node $i \in \mathcal{V}$.

The redundant node $i \in \mathcal{V}$ refers to element x_i of time series $x = (x_1, \dots, x_m)$. By

$$x' = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_m)$$

we denote the time series obtained from x by removing element x_i . Let $H'_k = (G'_k, \lambda'_k)$ be the resulting labeled warping graph, where λ'_k denotes the labeling function obtained by restricting λ_k to the subset $\mathcal{V}' \sqcup \mathcal{W}_k$ for all $k \in [N]$. Then the labeled warping graphs H'_k represent warping paths $q^{(k)}$ that align time series x' with sample time series $x^{(k)}$. By construction and definition of the weight function ω , we find that $\omega(H'_k) \leq \omega(H_k)$. Since h is monotonously increasing, we obtain

$$F(x) = \frac{1}{N} \sum_{k=1}^N h(\omega(H_k)) \leq \frac{1}{N} \sum_{k=1}^N h(\omega(H'_k)) = F(x').$$

By construction, we have $\ell(x') = \ell(x) - 1$. This completes the proof. \square

Proof of Corollary 2.8

Corollary. *Let $\mathcal{X} \in \mathcal{T}^N$ be a sample and let $\rho \in \mathbb{N}$ be the reduction bound of \mathcal{X} . Suppose that $\mathcal{F}_m \neq \emptyset$ for every $m \in [\rho]$. Then \mathcal{X} has a sample mean.*

Proof. For every $m \in [\rho]$ let F_m^* denote the restricted sample variation. We assume that $\mathcal{F} = \emptyset$. Then there is a time series $x \in \mathcal{T}$ of length $\ell(x) = p$ such that $F(x) = F_p(x) < F_m^*$ for all $m \in [\rho]$. This implies $p > \rho$, because otherwise we obtain the contradiction that $F_p(x) < F_p^*$. Let $q = p - \rho(\mathcal{X})$. By applying Theorem 2.7 exactly q -times, we obtain a time series $x' \in \mathcal{T}$ of length $\ell(x') = \rho$ such that $F_\rho^* \leq F(x') \leq F(x)$. This contradicts our assumption that $F(x) < F_\rho^*$. Hence, \mathcal{F} is non-empty. \square

Proof of Proposition 2.9

Proposition. *Let $\mathcal{X} \in \mathcal{T}^N$ be a sample. Suppose that \mathcal{A} is a finite attribute set. Then the following statements hold:*

1. $\mathcal{F}_m \neq \emptyset$ for every $m \in \mathbb{N}$.
2. $\mathcal{F} \neq \emptyset$.

Proof. Let $m \in \mathbb{N}$ be arbitrary. Since \mathcal{A} is finite, the set subset \mathcal{T}_m is also finite and consists of $m^{|\mathcal{A}|}$ time series. Then the set $F(\mathcal{T}^N)$ is a finite set. Hence, the restricted sample mean set \mathcal{F}_m is non-empty and finite. Since m was chosen arbitrarily, the first assertion follows. The second assertion follows from Corollary 2.8. \square

Proof of Proposition 2.10

Proposition. Let $\mathcal{X} \in \mathcal{T}^N$ be a sample. Suppose that the following assumptions hold:

1. (\mathcal{A}, d) is a metric space of the form $(\mathbb{R}^d, \|\cdot\|)$, where $\|\cdot\|$ is a norm on \mathbb{R}^d .
2. The loss h is continuous and strictly monotonously increasing.

Then the following statements hold:

1. $\mathcal{F}_m \neq \emptyset$ for every $m \in \mathbb{N}$.
2. $\mathcal{F} \neq \emptyset$.

The proof of Prop. 2.10 uses the notion of coercive function. A continuous function $f : \mathbb{R}^q \rightarrow \mathbb{R}$ is coercive if

$$\lim_{\|x\| \rightarrow \infty} f(x) = +\infty,$$

where $\|\cdot\|$ is a norm on \mathbb{R}^q .

Proof. We first consider the Euclidean norm $\|\cdot\|_2$ on some real-valued vector space \mathbb{R}^q . The Euclidean norm is coercive. Since h is continuous and strictly monotonously increasing on $\mathbb{R}_{\geq 0}$, the composition $h(\|x\|_2)$ is coercive and continuous. Every norm $\|\cdot\|$ on \mathbb{R}^q is equivalent to the Euclidean norm. Therefore, we can find constants $0 < c \leq C$ such that

$$c\|x\|_2 \leq \|x\| \leq C\|x\|_2$$

for all $x \in \mathbb{R}^q$. Hence, $h(\|x\|)$ is coercive and continuous for every norm on \mathbb{R}^q .

Suppose that $\mathcal{X} = (x^{(1)}, \dots, x^{(N)}) \in \mathcal{T}^N$ is a proper sample of N time series $x^{(k)}$ of length $\ell(x^{(k)}) = n_k \geq 2$ for all $k \in [N]$. Let $m \in \mathbb{N}$ be arbitrary. Expanding the definition of the restricted Fréchet function F_m gives

$$F_m(x) = \frac{1}{N} \sum_{k=1}^N h\left(\delta\left(x, x^{(k)}\right)\right) = \frac{1}{N} \sum_{k=1}^N \min \left\{ h\left(c_p\left(x, x^{(k)}\right)\right) : p \in \mathcal{P} \right\},$$

where $c_p(x, y)$ is the cost of aligning time series x and y along warping path p . Since $\mathcal{T}_m = \mathcal{A}^m = \mathbb{R}^{d \times m} = \mathbb{R}^q$, we can define the function

$$g_{p^{(k)}} : \mathbb{R}^q \rightarrow \mathbb{R}, \quad x \mapsto c_{p^{(k)}}\left(x, x^{(k)}\right) = \sum_{l=1}^{L_k} h\left(\left\|x_{i_l} - x_{j_l}^{(k)}\right\|\right),$$

where $p^{(k)} \in \mathcal{P}_{m, n_k}$ is a warping path with L_k elements aligning x and $x^{(k)}$. The function $g_{p^{(k)}}$ is continuous and coercive as a sum of non-negative continuous and coercive functions. Then $g_{p^{(k)}}$ has a global minimum.

We define the set $\mathcal{P}_m = \mathcal{P}_{m, n_1} \times \dots \times \mathcal{P}_{m, n_N}$. Then every element of \mathcal{P}_m is of the form $\mathcal{C} = (p^{(1)}, \dots, p^{(N)})$, where $p^{(k)}$ is associated to time series $x^{(k)}$ for all $k \in [N]$. Then we can equivalently rewrite the restricted Fréchet function $F_m(x)$ as

$$F_m(x) = \min \{F_{\mathcal{C}}(x) : \mathcal{C} \in \mathcal{P}_m\},$$

where the component functions $F_{\mathcal{C}} : \mathbb{R}^{d \times m} \rightarrow \mathbb{R}$ are functions of the form

$$F_{\mathcal{C}}(x) = \frac{1}{N} \sum_{k=1}^N g_{p^{(k)}}(x).$$

This shows that $F_{\mathcal{C}}(x)$ has a minimum. Let $F_{\mathcal{C}}^*$ denote the minimum value of $F_{\mathcal{C}}(x)$. From

$$\min_x F_m(x) = \min_x \min_{\mathcal{C}} F_{\mathcal{C}}(x) = \min_{\mathcal{C}} \min_x F_{\mathcal{C}}(x)$$

follows

$$\min_x F_m(x) = \min_{\mathcal{C} \in \mathcal{P}_m} F_{\mathcal{C}}^*.$$

Since \mathcal{P}_m is a finite set, we obtain that F_m has a minimum. This shows the first assertion, because m was arbitrary. The second assertion follows from Corollary 2.8. \square

4 Conclusion

The key contribution of this article is the Reduction Theorem. This theorem states that time series whose lengths exceed the reduction bound can be reduced to shorter time series without increasing the value of the Fréchet function by removing redundant elements. This result is useful, because it considerably simplifies existence proofs of a sample mean in DTW-spaces. The proof of the Reduction Theorem is framed into the theory of warping graphs. The second contribution presents sufficient conditions of existence of a sample mean in DTW-spaces. The proposed conditions are sufficiently general to cover customary variations of DTW-spaces. The proposed results theoretically justify existing pattern recognition applications of the sample mean in DTW-spaces, serves as a first step towards a statistical theory of DTW-spaces, and points to a technique that could help to improve candidate solutions of mean-algorithms.

Future research aims at exploring how the solution quality of state-of-the-art algorithms depends on the length of the argument of the Fréchet function, how redundant elements are distributed and to which extent their removal results in improved solutions. A second strand of research aims at developing a statistical theory of DTW-spaces that includes consistency results of the sample mean as next step.

References

- [1] W.H. Abdulla, D. Chow, and G. Sin. Cross-words reference template for DTW-based speech recognition systems. *Conference on Convergent Technologies for Asia-Pacific Region*, 2003.
- [2] A. Bhattacharya and R. Bhattacharya. *Nonparametric Inference on Manifolds with Applications to Shape Spaces*. Cambridge University Press, 2012.
- [3] I.L. Dryden and K.V. Mardia. *Statistical shape analysis*, Wiley, 1998.
- [4] A. Feragen, P. Lo, M. De Bruijne, M. Nielsen, and F. Lauze. Toward a theory of statistical tree-shape analysis. *IEEE Transaction of Pattern Analysis and Machine Intelligence*, 35:2008–2021, 2013.
- [5] M. Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l'institut Henri Poincaré*, 215–310, 1948.
- [6] C.E. Ginestet. Strong Consistency of Fréchet Sample Mean Sets for Graph-Valued Random Variables. arXiv: 1204.3183, 2012.
- [7] L. Gupta, D. Molfese, R. Tammana, and P.G. Simos. Nonlinear alignment and averaging for estimating the evoked potential. *IEEE Transactions on Biomedical Engineering*, 43(4):348–356, 1996.
- [8] V. Hautamaki, P. Nykanen, P. Franti. Time-series clustering by approximate prototypes. *International Conference on Pattern Recognition*, 2008.
- [9] S. Huckemann, T. Hotz, and A. Munk. Intrinsic shape analysis: Geodesic PCA for Riemannian manifolds modulo isometric Lie group actions. *Statistica Sinica*, 20:1–100, 2010.
- [10] B.J. Jain. Statistical Analysis of Graphs. *Pattern Recognition*, 60:802–812, 2016.
- [11] D.G. Kendall. Shape manifolds, procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*, 16:81–121, 1984.
- [12] P.T. Kim and J.-Y. Koo. Comment on [9]. *Statistica Sinica*, 20:72–76, 2010.
- [13] J.B. Kruskal and M. Liberman. The symmetric time-warping problem: From continuous to discrete. *Time warps, string edits and macromolecules: The theory and practice of sequence comparison*, pp. 125–161, 1983.
- [14] J.S. Marron and A.M. Alonso. Overview of object oriented data analysis. *Biometrical Journal*, 56(5):732–753, 2014.
- [15] V. Niennattrakul and C.A. Ratanamahatana. Shape averaging under time warping. *IEEE International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*, 2009.
- [16] F. Petitjean, A. Ketterlin, and P. Gancarski. A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition* 44(3):678–693, 2011.
- [17] L.R. Rabiner and J.G. Wilpon. Considerations in applying clustering techniques to speaker-independent word recognition. *The Journal of the Acoustical Society of America*, 66(3): 663–673, 1979.
- [18] H. Sakoe and S. Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, 1978.
- [19] S. Soheily-Khah, A. Douzal-Chouakria, and E. Gaussier. Generalized k-means-based clustering for temporal data under weighted and kernel time warp. *Pattern Recognition Letters*, 75:63–69, 2016.
- [20] H. Wang and J.S. Marron. Object oriented data analysis: sets of trees. *The Annals of Statistics* 35:1849–1873, 2007.